ISCR

Institute for Scientific Computing Research

University Collaborative Research Program Subcontract Research Summaries

CASC Center for Applied Scientific Computing

Software Infrastructure for Multi-Tier Implementation of Structured Adaptive Mesh Hierarchies

Scott Baden
University of California, San Diego

Summary:

We have continued our investigation into scalable parallel programming methodology and performance tradeoffs arising in mixed-mode (message-passing, multi-threaded) programming of hierarchically constructed multi-tier computers. We have leveraged previous work with the KeLP system for manipulating hierarchical data structures, which is currently running on ASCI Blue Pacific. Our investigation targets structured adaptive mesh refinement, including hierarchical load balancing strategies. A uniformly refined multi-block code, including a load balancer, was developed first, followed by a structured adaptive mesh version. An existing implementation of an SAMR Poisson solver, written in a single-tier dialect of KeLP, was ported to hierarchical KeLP. The investigation will deliver a computational testbed permitting LLNL scientists to explore performance tradeoffs under mixed-mode parallelism for structured adaptive mesh refinement and a variety of other block-structured applications of interest to the laboratory. We are conducting comparative studies with the NPACI Tflops architecture, which is based on Power3 SMPs of an architecture similar to Livermore's ASCI White.

Principal LLNL contact: John May, CASC

High Resolution Parallelized Self-Organizing Maps

Jackson Beatty
University of California, Los Angeles

Summary:

This project was designed to create publicly available parallel algorithms for computing self-organizing neural networks. However, it also gave rise to two major new programs at UCLA. The first is a research program seeking to establish the computational basis of cortical information processing in the human brain. The second is a program of instruction in neuroinformatics at UCLA, funded by the Human Brain Project, which is the first such program in the United States.

Four new parallel algorithms for computing and extending Kohonen's self-organizing map (SOM) algorithm were created under this grant. The first was psom.fortran, a generalized version of the SOM algorithm coded in Fortran and utilizing the MPI library for parallel computation. The psom.fortran algorithm freed the Kohonen procedure from several of its arbitrary parameters and provided a general mechanism for understanding unsupervised competitive learning processes embedded within a two-dimensional surface, similar to that of the human cerebral cortex. The second algorithm was psom.c, which is algorithmically identical to psom.fortran but was completely rewritten because of MPI library difficulties with Fortran on Livermore Teracluster. The third algorithm we call Adaptasom, for adaptive coordinate visualization of self-organizing map formation. Adaptasom accepts information from either psom.c or psom.fortran as a map is being generated and displays it as movements of reference vectors in the self-organizing map space. Adaptasom is based on insights and algorithms developed by Rauber Andreas in his doctoral thesis at the Technical University of Vienna in 1996. The fourth algorithm is called Retina, which is a set of procedures for transforming calibrated digital photographs of natural images into a representation similar to that provided by the retina to the optic nerve and hence to human visual cortex. Retina provides the primary means of testing the Neuronal Empiricism Hypothesis and for discovering the common computational algorithm of the primate and human cerebral cortex.

For decades neuroscientists have studied the properties of cells in various specialized areas of the cerebral cortex and detailed their unique properties. Particularly well studied are the primary sensory areas, those cortical regions that receive the initial input from the different senses. Thus, the properties and organization of cells in the primary visual cortex are by now well known. Most of these cells respond best to lines of a particular orientation in a specific region of the retina. These cells form a part of a pinwheel, the segments of which contain cells favoring a specific particular organization. The pinwheels themselves are organized into a spatial mapping of the retina, so that each particular region of the cortex exclusively processes information from a unique portion of the retina.

Thus, it was of profound importance to learn that an identical pattern of neural organization emerges in auditory cortex when visual information is surgically redirected from visual to auditory cortex in newborn ferrets. The only possible explanation for this finding is that the cellular properties of visual cortex neurons are not genetically determined (as is usually assumed) but rather are extracted from information provided by the visual system using a computational algorithm that is common to all cerebral cortex. This is an extremely powerful idea, which forms the basis for what I term the Neuronal Empiricism Hypothesis. In its simplest

Summary (continued):

form, it simply states that the cortex learns to extract regular features from whatever sensory information it is provided, because such regularities provide information about real objects in the organism's environment.

The striking pattern of results that we are beginning to observe from our psom algorithm operating on the output of retina processing calibrated photographs of natural scenes suggests that the Kohonen algorithm (or some similar algorithm) may be capable of generating the observed properties and organization of cortical visual neurons simply by extracting regularities from the sensory stream provided by the retina. This finding, should it prove robust, would constitute a major advance in the field of computational neuroscience and provide deep evidence that there in fact exists a common algorithm for cortical computation.

This heavily computational but data-constrained approach to understanding the human nervous system also has educational implications. Our LLNL grant has provided us at UCLA with sufficient experience in large-scale computation to formulate a new program in neuroscience instruction at UCLA, called neuroinformatics.

One grand challenge facing science today is to transform neuroscientific data into facts and from these facts extract knowledge and understanding. Neuroinformatics offers a set of powerful tools by which this transformation of data may be accomplished. The new neuroinformatic tools include the creation of systematic, content-driven databases, sophisticated data visualization methods, unsupervised data-driven pattern extraction methods, neural modeling across a wide range of levels, and electronic collaboration. Similar methods have been applied with striking success to the (simpler) problems of gene and protein sequencing, an enterprise that created the discipline of bioinformatics. All of these new tools are examples of new information technology and—taken together—they have redefined the field of scientific computing.

Thus, neuroinformatics brings together the methods and concepts of neuroscience, informatics, and computation that are necessary to understand the human brain in all its complexity. Although computers have played important roles in neuroscience research for decades, it has only been in the last several years that the economics of computer manufacturing and marketing have made large-scale supercomputing technology available for the study of the brain and its functions. Neuroinformatics offers neuroscience new opportunities for extracting understanding from data and meaning from measurements.

With the support of the Human Brain Project, we are beginning to prepare a new generation of informatically oriented neuroscientists by specifically introducing them to neuroinformatic methods and principles while they are still undergraduates, and to make the curricula by which we teach neuroinformatics to our students freely available to colleges and universities throughout the world. This instructional program grew directly from the computational studies that we carried out over the last two years at LLNL.

Principal LLNL contact: Farid Dowla, Electronics Engineering Technology

Semiclassical Quantum-Mechanical Modeling using Parallel PIC Methods

John Dawson and Viktor Decyk

University of California, Los Angeles

Summary:

We have been successful in modeling many particle quantum systems by combining a semiclassical approximation of Feynman path integrals with parallel computing techniques previously developed at UCLA for simulating plasmas. Issues regarding the derivation and implementation of these techniques are described briefly, a method to address quantum particle statistics is introduced, and studies of two quantum-mechanical problems—a two-electron, one-dimensional atom and electrostatic quasi-modes due to quantum effects in a hot electron plasma—are discussed.

We have continued the collaborative program with Denis Hewett at LLNL on the use of parallel computers for quantum physics. The effort concentrates on building models and developing techniques to simulate systems of practical interest. Our focus is on multiparticle systems where quantum effects are significant.

UCLA has pioneered particle methods (PIC) for modeling plasmas. For over six years, we have developed efficient and effective methods for using parallel computers for carrying out PIC simulations (90% parallel efficiency, 40% of estimated peak speed). At present we are able to handle over 200 million particles in three dimensions using self-consistent full electromagnetic fields. Our methods are robust and portable and have run successfully on a wide range of computers (e.g., Cray-90s, T3Ds, T3Es, and SP2s).

The focus of the project has been to apply PIC techniques to the semiclassical approach and develop a dynamical model of many particle quantum systems. This model would allow detailed investigation of optical properties, ionization potential, conductance, and a host of other experimentally determined material properties; it has predictive capabilities useful for the design and understanding of devices where quantum effects are important. The results have been very promising. This work is the part of the Ph. D. thesis of Dean E. Dauger.

We have made good progress towards a dynamic quantum mechanical modeler and have accomplished two major milestones: we have successfully transformed a plasma code into a quantum code, and we have applied this code to a range of quantum problems, including multiple particles.

The theoretical basis for the methods developed here is the Feynman path integral, which can predict the time evolution of a wavefunction using a sum over all possible paths. However, based on Van Vleck's Green's function, the path integral may be reduced to a sum over all possible classical paths, using the semiclassical approximation. This approach is easily extended to multiple dimensions and to include full electromagnetism. It turns out the primary cost of this method is in tracing the classical trajectories through electromagnetic fields.

Summary (continued):

We considered how to handle interactions between multiple particles. For typical interactions described in a multiparticle Hamiltonian, we know how to derive an expansion for the effective Hamiltonian for any one particle. The ordering coefficient of this expansion is proportional to the time of the wavefunction push, implying a recommended limit for how long to evolve the wavefunction before reconstructing the wavefunction and updating the fields. Spin-spin interactions can be incorporated very easily into this structure.

Our investigation has also focused on the accuracy of the techniques essential to this project. We have produced and examined a precise derivation of the semiclassical method as it applies to our problem and have addressed a significant issue in using semiclassical methods. When time evolving one quantum particle over one quantum time step, the derivation describes that the sum over all paths between an initial and a final point simplifies to one classical path. This derivation also produces a coefficient, which leads to caustics and related potential difficulties in using semiclassical methods. Classical path times in previous use of semiclassical methods were long enough for caustics to pose a significant challenge. However, because our approach repeats many quantum time steps to evolve the wavefunctions, path times in our use of these methods are much shorter, making caustics and related concerns unlikely in the systems with which we are interested.

Another issue we investigated was the application of these techniques to systems exhibiting fermion statistics. This issue is of importance when applying these techniques to the simulation of electrons.

The first system we applied this technique to was two fermions in an infinite square well. We chose this potential system because it had a well-known single particle energy spectrum yet did not have the regularity of the simple harmonic oscillator system. Peaks in the energy spectra were found where we thought they should be, at energies that correspond to two-particle states of differing quantum number. More importantly, no peaks were found where we thought they should not be, that is, at those energies that would contradict the Pauli exclusion principle and energies in between expected energies.

We investigated further fermionic multiparticle problems. The next system was a one-dimensional atom: two mutually repulsive fermionic electrons attracted to a central, positive, one-dimensional charge. This system was chosen for its appropriateness to the code while being significantly more difficult to solve analytically. As with the infinite square well case, we found the energy spectra in the one-electron case and the two-electron case, and we found the ground-state energy to be comparable to what we could predict analytically using variational methods.

Summary (continued):

There is a great deal of what we can do with the code, which we have used for problems with over a hundred quantum particles. We intend to pursue further analysis of the quantum fluctuations problem, perhaps extending it to predict solar reaction rates. Also, it is possible to investigate particular issues regarding the implementation. One issue is to resolve why the boundary conditions for the infinite square well require the virtual classical particles bounce one half grid beyond the boundary. Another is to consider alternative decompositions for better parallelism, as the current decomposition over space is not ideal when the number of grids is low. Third, the current virtual particle sampling scheme blankets phase space. There may be ways of simplifying or optimizing this distribution to reduce the computational cost spent to evolve a wavefunction.

Further in the future, a number of additional features should offer a more complete modeling of quantum mechanics. Other higher order effects can be included in the code. Techniques at implementing two and three dimensions and full electromagnetism should carry over well from the plasma PIC codes. Along with higher dimensions, we can fully incorporate spin phenomena. With these additions, the code would be able to model multielectron atoms, chemical reactions, quantum electronics, and solid state physics, among a host of addressable physical problems.

We are applying the code to a problem concerning fluctuations in the electrostatic spectrum in a hot electron plasma due to quantum mechanics. A particular analysis of this problem may have important consequences for astrophysical predictions about stellar evolution and the early universe. It has been proposed that the energy density in the electromagnetic spectrum of a hot plasma when including quantum effects is significantly different than what is predicted using classical theory alone, which we intend to test.

Principal LLNL contact: Dennis Hewett, CASC

Development of Bioinformatic Tools for the Integration and Analysis of Microarray Data

Jeffrey Gregg
University of California, Davis

Summary:

DNA microarrays offer the ability to determine gene expression values of hundreds to thousands of genes simultaneously, offering insight into the complex behavior of the cells being studied. However, there currently exists no robust system for cataloguing and analyzing the massive amounts of data being generated in these micorarray studies. Our principal objective has been to develop, evaluate, and implement a flexible database schema and develop large-scale computational methodologies for the analyses of expression microarray data generated by the projects already initiated in Dr. Gregg's and Dr. Wyrobek's research laboratories at UC Davis and LLNL, respectively. To this end, we are designing and implementing appropriate database structures for support of gene expression data at UC Davis and LLNL and we are developing analytical tools for the interpretation of microarray data and its correlation with cellular biomarker measurements and clinical outcomes. Development of an effective database management system and analytical tools will allow interpretation of biologically relevant gene expression data that will impact the understanding of human disease.

Principal LLNL contact: Andrew Wyrobek, Biology and Biotechnology

Tailor-made Virtual Libraries

B. S. Manjunath
University of California, Santa Barbara

Summary:

Different people can differ significantly on the similarity of two multimedia objects, such as still images or videos. To an urban planner, two regions are similar because of the distribution of buildings, parks, and roads. However, when an analyst in a forest department looks at these images of urban regions adjacent to forests, her discrimination could be based on a different set of attributes, including factors such as a possible fire hazard because of dead wood content. To a botanist two pictures can be similar because they contain views of the same flower, while to a casual person the images might be dissimilar because the backgrounds are quite different. It is difficult, in general, to extract image and video features that support all possible user perspectives. This in turn makes it hard to access such multimedia databases by content.

We have outlined a strategy that could lead to the efficient access of large multimedia databases that is tailored to individual perceptual criteria. Our method is capable of learning the user's perceptual metric and using it to access the database by content. By learning these personal views and making them available in a useful manner, the knowledge and information content of the database grows. The database is no longer just a repository of bits and bytes, but also of knowledge and information that grows dynamically and is available in a concrete useful manner. Over time a layer of meta-data arises over the database providing multiple views of what was once a monolithic object. Thus, our objective is to develop a strategy for efficient access of large multimedia databases that is tailored to individual perceptual criteria.

The focus of this research matches quite well with the ISCR focus area of large-scale data mining and pattern recognition. We collaborate with Imola Fodor, Center for Applied Scientific Computing at LLNL, on this research. Dr. Fodor's group is developing scalable algorithms for the interactive exploration of large, complex, multi-dimensional scientific data, and the UCSB group's work on integrating dimensionality scaling and learning perceptual metric directly contributes to this effort. The techniques developed in this project have a wide range of applications. To compute the mapping from the image features to the perceptual space, new tools for dimensionality reduction have been developed. The mapping from the user's perceptual criteria to a low dimensional vector space is useful in developing fast indexing and retrieval strategies in large scientific databases, and in data mining.

We are building a software prototype for large scientific image databases (such as satellite or aerial imagery). The prototype will be used to test various techniques for the three components, i.e., creation of feature vectors (XVs), computing the c-Maps, and adaptive indexing.

Principal LLNL contact: Imola Fodor, CASC

Development of a Three-Dimensional Relativistic Particle-in-Cell Code for Studying the Production of Useful Electron Bunches Using Ultra-Intense Laser Pulses

Warren B. Mori
University of California, Los Angeles

Summary:

We have developed a parallelized three-dimensional relativistic particle-in-cell (PIC) code for studying the production of useful electron bunches using ultra-intense laser pulses. The electron bunches are produced when the radiation pressure of an intense laser either directly accelerates electrons or excites a plasma wave wake which accelerates electrons. Ultra-short electron bunches could be useful for injection into high-energy particle accelerators, for radiation sources, and for the fast ignitor fusion concept. A reliable and robust code will be invaluable towards the realization of such ultra-short electron bunches in a laboratory.

Our object-oriented Fortran90 parallelized PIC code has been used to model intense electron and laser pulses propagating through tenuous plasmas. Based on large-scale runs, we have carried out performance optimizations. We plan to add additional current/charge deposition schemes and to continue to work with Dr. Langdon's group to benchmark various 3D parallel algorithms.

Principal LLNL contact: Bruce Langdon, X Division

Sensitivity and Uncertainty Analysis for Large-Scale Differential-Algebraic Systems

Linda Petzold
University of California, Santa Barbara

Summary:

Sensitivity analysis of large-scale differential-algebraic (DAE) systems is important in many engineering and scientific applications. Sensitivity analysis generates essential information for design optimization, parameter estimation, optimal control, model reduction, process sensitivity and experimental design. Our recent work on methods and software for sensitivity analysis of DAE systems has demonstrated that forward sensitivities can be computed reliably and efficiently via automatic differentiation in combination with DAE solution techniques designed to exploit the structure of the sensitivity system. The DASPK3.0 software package was developed for forward sensitivity analysis of DAE systems with index up to two, and it has been used in sensitivity analysis and design optimization of several large-scale engineering problems.

Some problems require the sensitivities with respect to a large number of parameters. For these problems, particularly if the number of state variables is also large, the forward sensitivity approach is intractable. In the past year we have begun an investigation of adjoint (reverse) sensitivity methods, which have the potential to overcome this limitation. We have derived the adjoint system for general DAEs and investigated some of its fundamental properties. In particular, we have derived the adjoint DAE systems by a variational method, along with conditions for the consistent initialization of index-1 and Hessenberg index-2 adjoint DAEs. We have investigated the stability of the adjoint DAE system and shown that for semi-explicit DAE systems the adjoint DAE is stable whenever the original DAE is stable. For fully implicit DAE systems this is no longer true, but we can show that the stability is preserved by an augmented form of the adjoint system. Thus, it is this system that we discretize for the numerical solution.

Work has begun and is now well underway on an adjoint DAE solver, DASPK-ADJOINT, which uses a modified DASPK3.0 to solve the forward and reverse problems. Some preliminary work has been done on characterizing the class of well-posed sensitivity problems for DAEs, and on analysis of the numerical stability for fixed stepsizes of the BDF methods applied to the augmented adjoint DAE. Complexity analysis and preliminary numerical results indicate that the adjoint sensitivity method is advantageous over the forward sensitivity method for applications with a large number of sensitivity parameters and few derived functions.

During the course of this work, the PI visited LLNL several times and gave a lecture at the LLNL CASC Nonlinear Solvers Workshop in July 2000.

Principal LLNL contact: Steven Lee, CASC

Global Simulation of the Earth's Magnetosphere with Adaptive Mesh Refinement

Joachim Raeder
University of California, Los Angeles

Summary:

We have added structured adaptive mesh refinement capabilities to the UCLA Global Geospace Circulation Model (UCLA-CGCM) in collaboration with the SAMRAI project team in the Center for Applied Scientific at LLNL. The UCLA–GGCM is a global circulation model of Earth's magnetosphere and ionosphere that has been used for several years to study the interactions of the solar wind, magnetosphere, and ionosphere. The UCLA–GQCM with local adaptive mesh refinement supports a grid resolution in critical regions that is about two orders of magnitude better than the base code. This additional resolution allows us to address problems that were previously elusive to global modeling, such as flux transfer events, the formation of boundary layers, and the formation of thin current sheets in the late substorm growth phase.

SAMRAI is a general software support framework for structured adaptive mesh refinement (AMR) applications on parallel high performance computing hardware. In collaboration with the SAMRAI group, we will continue to implement a new GGCM code that integrates the existing numerical kernels from our current application into the SAMRAI framework. We will in particular address the multiphysics integration of the ionosphere submodel, and issues of visualization and the efficiency of the time integration algorithm in the adaptive mesh context.

Principal LLNL contacts: Richard Hornung and Scott Kohn, CASC

Theoretical Foundations of Single Bubble Sonoluminescence

Andrew Szeri
University of California at Berkeley

Summary:

Sonoluminescence (SL) is the production of light from micro-bubbles driven into extreme volume oscillations by a strong acoustic field. The temperatures and pressures in current experimental efforts are intense enough to produce, momentarily within the bubble, a cool dense plasma which emits light like a tiny star from the surface of an opaque core. The size of the light emitting region, and hence the luminosity of the bubble, depend sensitively on the physical characteristics of the material inside the bubble, and on the forcing to which it is subjected. Significant progress has been made on the theoretical understanding of SL, including the consequences of material selection, long-term chemical effects on the physics, and modeling techniques. We developed an accurate but fast computational model of heat and mass transfer in SL bubbles, based on a more detailed but very time-consuming model. The reduced model was subsequently applied to the study of different choices of materials (cryogenic noble gas systems, exotic low-vapor-pressure liquids) and the long-term consequences of chemical activity (the dissociation hypothesis).

Just prior to the project, we had written a detailed simulation to elucidate the role of water vapor in SL. We found that significant amounts of water vapor enter the bubble during the long slow expansion, as other researchers have shown, but that much water vapor is incapable of diffusing to the bubble wall in time to exit the bubble at the collapse. At higher and higher amplitudes of the driving, more and more water vapor is trapped in the bubble at the collapse. This water vapor is involved in chemical reactions, which are accelerated by the compression heating and high density of the bubble. Because the reactions are endothermic, they reduce the heating of the bubble, as does the lowering of the ratio of specific heats of the gas mixture. This is strongly coupled to observables: the number of photons and pulse width of light emitted. That study gave us the foundation we needed to develop a much faster model of reduced complexity. In this model, we take account of heat and mass transfer, as well as chemical reactions, but the fields of interest (pressure, temperature, and composition) have been averaged across the volume of the bubble contents. This leads to ordinary differential rather than partial differential equations. Despite the *simplicity* of the reduced model, comparison of results with the detailed model showed excellent agreement in terms of temperatures achieved, OH radical production, water vapor trapped, etc.

Summary (continued):

During the course of the research, in collaboration with William Moss of LLNL, we realized that the usual corrections for "compressibility" to the Rayleigh–Plesset equation for bubble dynamics had missed an important first-order effect. The usual corrections account for *liquid* compressibility, i.e., sound waves in the liquid, which contribute to the damping of the bubble dynamics by the radiation of acoustic energy. We realized that of even greater importance was compressibility of the gas, i.e., development of a nonuniform pressure field during rapid collapse or expansion of the bubble. We developed a simple correction to account for this effect, which yields much better agreement with experimental results for strongly forced (SL) bubbles.

We have also made use of the reduced model in the study of cryogenic systems. There was speculation that, for example, a Helium bubble in liquid Argon might be an ideal system for SL. The reason is that there would be no reduction of the internal temperature due to thermodynamic or chemical processes. Unfortunately we found that such bubbles would likely be overwhelmed by the large quantity of vapor that flows in during the expansion. We studied in addition other non-aqueous systems. These are challenging computations, as a multi-dimensional parameter space must be investigated in order to determine the narrow range of parameters where SL can be achieved.

As a final application of the reduced model we are studying the dissociation hypothesis of Lohse and coworkers. The idea is that diatomic components of air in an SL bubble undergo chemical reactions, resulting in soluble products that leave the bubble nearly pure Argon. This happens over several seconds, it is thought, which corresponds to perhaps 50–60,000 complete bubble oscillations. This places extreme demands on our computational model. We have developed the first simulation capable of computing this effect, but we must speed up the calculation further to develop data for comparison with experiments over a range of parameters.

Principal LLNL contact: William C. Moss, Geophysics and Global Security